

Spur-Reduced Digital Sinusoid Generation Using Higher-order Phase Dithering

M. J. Flanagan

Electrical Engineering Department
and Jet Propulsion Laboratory
California Institute of Technology
Pasadena, CA 91109

G. A. Zimmerman

Jet Propulsion Laboratory
and Electrical Engineering Department
California Institute of Technology
Pasadena, CA 91109

Abstract

A higher-order phase dithering technique that reduces the complexity of digital sinusoid generation is presented and analyzed. M^{th} order phase dithering denotes the addition of M independent, uniform variates to the sinusoid phase prior to word-length reduction. Spurious magnitudes due to quantization effects are accelerated from the usual -6 dBc per phase bit to $-6(M+1) \text{ dBc}$ per phase bit, while the noise power increases linearly in M . For a given spurious specification, higher-order phase dithering permits the use of fewer phase bits than the non-dithered system, resulting in an exponential reduction in look-up table size, dramatically decreasing system complexity.

1 Introduction

Phase dithering is a useful way to reduce the complexity of digital sinusoid generators [1]. The addition of dither ([2], [3]) prior to a reduction in the phase word length permits the retention of fewer phase bits to maintain a given spurious specification. Since the phase is used to specify an address in systems using an amplitude look-up table, the use of fewer phase bits results in an exponential reduction in memory. This can provide substantial savings in VLSI implementations. Complexity is also greatly reduced in generators that use algorithms instead of look-up tables to produce output amplitude samples.

Phase-dithered digital sinusoid generators in digital signal processing applications produce the same effects as frequency synthesizers with phase noise in analog applications. While the savings in VLSI complexity are obtained at the expense of increased system noise, overall degradation is modest in high-bandwidth systems where complexity reductions are most keenly appreciated. The increased system noise is also a minor concern in applications that process all but the highest signal-to-noise ratio inputs.

This paper presents a new approach to phase dithering. The dither signal is constructed by summing M independent sequences of i.i.d. variates. Each variate is uniformly distributed over one quantization interval, e.g., $[0, A)$ or $[-A/2, A/2)$, where $A = 2\pi/b$, and b is the number of bits used to represent the phase

after word-length reduction. This signal is referred to as an M^{th} order dither signal because its characteristic function has M^{th} order zeros at non-zero integer multiples of $\frac{2\pi}{A}$. This approach includes the method of [1] as a special case where $M = 1$ and also treats higher-order dithering where $M > 1$.

2 General Analysis

The overall noise and the spurious performance are the two figures of merit in a phase-dithered digital sinusoid generator that are considered here. Amplitude quantization effects are ignored by assuming a sufficiently large number of bits are used to represent each amplitude sample. Work has been done treating both amplitude and phase dithering [4] where it has been noted that phase dithering provides greater complexity savings than amplitude dithering since the look-up table size grows exponentially in the number of phase bits and only linearly in the number of amplitude bits. In practice, dither signal constraints can be relaxed to allow for simple generator structures amenable to VLSI implementation. Such structures are presented in a later section.

It is desired to find the variance of the noise introduced by M^{th} order phase dithering. Let $x[n] = \cos(2\pi\phi[n])$ represent a digital sinusoid that has been generated by a system using M^{th} order phase dithering. Note that amplitude quantization effects are being ignored. The normalized phase signal is written as $\phi[n] = \phi_0[n] + \epsilon[n]$, where $\phi_0[n]$ is the ideal phase signal and $\epsilon[n]$ is the quantization noise signal. The normalized phase sequence is the digital signal from which the system constructs the amplitude sample sequence, $x[n]$. The normalized phase signal is generated by quantizing the sum of the ideal phase signal and an M^{th} order dither signal, $z[n]$. The dither signal, $z[n]$, is an i.i.d. sequence constructed by summing M independent variates, each uniformly distributed over one quantization interval. The quantization interval is defined by the quantizer that the normalized phase sequence emerges from. This type of a quantizer system is called a "non-subtractive dither system" and was studied in detail by Gray and Stockham

[3] who showed that the quantization noise signal is white with variance $\frac{(M+1)\Delta^2}{12}$. Define the error due to phase dithering to be $e_x[n] = x[n] - \cos(2\pi\phi[n])$. Using small angle approximations and following the analysis of [1], the error due to phase dithering is easily shown to be a white process with variance $\frac{(M+1)\pi^2\Delta^2}{3}$. Since the noise power is white, it is often tolerable in wide-band systems.

The spurious performance of the phase dithered system is found by considering the autocorrelation function, $E\{x[n]x[n+m]\}$, where $E\{\cdot\}$ denotes the expectation operator and the lag, m , is non-zero. Since the lag is non-zero, $x[n]$ and $x[n+m]$ are independent because the dither signal, $\phi[n]$, is i.i.d. The autocorrelation function becomes the product of two expectations and it suffices to consider the "expected waveform," $E\{x[n]\}$, to obtain spectral information. Since the dither signal and the quantization error are $O(A)$ and A is small in practice, we can expand the output, $x[n]$, in a Taylor series about $e[n] = 0$:

$$x[n] = \sum_{k=0}^{\infty} \frac{(2\pi)^k e^k[n] f^{(k)}(2\pi\phi[n])}{k!}, \quad (1)$$

where $f^{(k)}(\theta)$ denotes the k^{th} derivative of the function $f(\theta) = \cos(\theta)$. The k^{th} term in the expansion is proportional to the quantization noise signal raised to the k^{th} power. To evaluate the expected waveform, we take the expected value of the Taylor series:

$$E\{x[n]\} = \cos(2\pi\phi[n]) + \sum_{k=1}^{\infty} \frac{(2\pi)^k E\{e^k[n]\} f^{(k)}(2\pi\phi[n])}{k!} \quad (2)$$

It was shown by Gray and Stockham [3] that the first M moments of the quantization noise, $e[n]$, are independent of the "desired" signal, $\phi[n]$. Therefore, each of the first M terms in the above sum over k can be written as a constant times either $\cos(2\pi\phi[n])$ or $\sin(2\pi\phi[n])$. Since each of these terms has frequency support only where the desired sinusoid, $\cos(2\pi\phi[n])$ has support, these terms do not generate spurious harmonics. The first M terms make the effective amplitude of the ideal portion of the output signal, $x[n]$, different from unity. The difference is small, however, because the largest of the first M terms is $O(A)$.

The $(M+1)^{\text{th}}$ moment of the quantization noise will generally be dependent on the ideal phase signal, $\phi[n]$, which is a periodic sawtooth signal. Therefore, the $(M+1)^{\text{th}}$ moment is generally periodic and gives rise to spurs which are scaled by $\frac{(2\pi)^{M+1}}{(M+1)!}$ and amplitude modulated by the sinusoid $f^{(M+1)}(2\pi\phi[n])$. Since $e[n]$ is $O(A)$, its $(M+1)^{\text{th}}$ moment is $O(\Delta^{M+1})$. Therefore, the maximum spur power is $O(\Delta^{2(M+1)})$, giving the $6(M+1)$ dBc per phase bit result. In order to tighten the bound on the maximum spur power, the following section considers the constant term in the order expression $O(\Delta^{2(M+1)})$.

3 Tight Performance Bounds for Specific-Order Dithering

While the theoretical bounds derived in the preceding section demonstrate the fundamental $-6(M+1)$ dBc per phase bit spurious performance for A^4 -order phase-dithered systems, comparison with the results for first-order dithering presented in [1] indicates that the actual spurious performance is significantly better. By examining the expected waveform as a power series in the quantization noise, and directly analyzing the second moment of the quantization noise, the worst-case spurious performance was found to be $\frac{\pi^4\Delta^4}{16}$, or $7.84 - 12.04$ dBc, for b-bit phase after first-order dithering.

Specific higher-order dither cases can be analyzed using the same techniques as those for first-order dithering. For second-order dithering, consider the third moment of the quantization noise, $E\{e^3[n]\}$. If the dither is the sum of two uniformly-distributed $[0, A)$ random variates, then it will have the following probability density function:

$$f(r) = \begin{cases} x/\Delta & \text{for } 0 \leq x < \Delta \\ (2 - x)/\Delta & \text{for } \Delta \leq x < 2\Delta \end{cases}$$

This dither variate is added to the phase value, and the result is truncated to form a b-bit value. For each phase sample value, $\phi[n]$, the quantization noise at a given time, $e[n]$, is a well-defined, discretely-distributed random variate. In second-order dithering, the quantization noise has three possible values:

$$e[n] = \begin{cases} -p[n]\Delta & \text{with probability } (1-p)^2 \\ (1-p[n])\Delta & \text{with probability } \frac{1}{2} - p + \frac{p^2}{2} \\ (2-p[n])\Delta & \text{with probability } \frac{p^2}{2} \end{cases} \quad (3)$$

where $p[n] = (\phi[n] \bmod A)/A$, the quantization error for non-dithered truncation normalized to $A = 1$.

A little algebra confirms that $E\{e[n]\}$ and $E\{e^2[n]\}$ are, as predicted by theory, independent of the input phase, as well as indicating that $\phi[n]$ for this truncation implementation should be biased to counteract the nonzero, constant value of $E\{e[n]\}$. Further algebra gives the following result for the third moment of the quantization noise:

$$E\{e^3[n]\} = \Delta^3 \left(-p[n]^3 + \frac{3p[n]^2}{2} - \frac{p[n]}{2} + \frac{1}{2} \right) \quad (4)$$

Over the valid range of $p[n]$, $0 \leq p[n] < 1$, the values of $E\{e^3[n]\}$ are tightly constrained so that $0.45A^3 < E\{e^3[n]\} < 0.55A^3$. This gives the peak-to-peak variation of any periodic component of the quantization noise sequence as $< 0.1A$. Following the analysis of [1] for first-order dithering, the second-order dithered system's spurious performance can be bounded more closely as:

$$SpSR \leq \left[\frac{(2\pi\Delta)^3 0.1}{6 \cdot 2} \right]^2 = 0.3 - 18.1 \text{ dBc} \quad (5)$$

where $SpSR$ is the ratio of the power in the maximum spur component to the power in the desired sinusoidal signal. It is not known whether this bound is achieved as it is in the first-order dithered system, but simulation examples shown in Section 5 indicate that it is tight within 1 dB.

Similar specific-order analysis can be performed for higher-order dithered systems, as design requirements are formed, after using the simpler $6(M+1)$ dBc per phase bit bound to give the approximate dither order required.

4 Performance and Complexity

The following comparison demonstrates the circuitry complexity savings of higher-order phase dithering over no phase dithering and first-order phase dithering. It is desired to have better than -90 dBc spurious performance when the clock rate is 100 MHz. This high level of spectral purity is desired for radio-science applications in NASA's Deep Space Network, as well as for use with high-resolution, high-dynamic range spectrum analyzers such as those in [5].

An often used and simple method of generating a sample of a sinusoid from a phase sample is to use a look-up table. The circuit complexity of a look-up table is exponential in the number of phase bits used, requiring 2^{b-2} entries, even when quadrant symmetries are employed. Achieving -90 dBc spur performance without phase dithering requires 17 bits of phase, as input to a look-up table, and often requires the use of arithmetic algorithms which involve multi-bit multipliers. By incorporating one additional adder and a pseudorandom number generator to the undithered system, a first-order phase-dithered sinusoid generator requires only 9 look-up phase bits for better than -90 dBc performance, a reduction in complexity of a factor of 256. Using two pseudorandom number generators and two additional adders over the undithered system, second-order phase dithering requires only 5 look-up phase bits to achieve -90 dBc performance. Section 5 will show that a simple, linear feedback shift register (LFSR) pseudorandom number generator will suffice. The resulting second-order phase dithering system memory consists of only 8 entries when quadrant symmetries are used, in contrast to the 32,768 entries that a system with no phase dithering would require, as shown in Table I. Also shown in the table is the penalty paid in the dithered systems: a white noise power term modulated onto the sinusoid which increases in power 6 dB with each reduced bit, and linearly with the order of the dither.

Table 1: System Comparison:

-90 dBc spurious and 100 MHz clock			
Phase Dither Order	Look-up Base Bits	Memory Entries	Noise Power Spectral Density
none	17	32768	-173.8 dBc/Hz
1	9	128	-124.2 dBc/Hz
2	5	8	-97.0 dBc/Hz

The noise power spectral density of the second-order phase-dithered system is higher than the phase

noise density of economically priced analog frequency synthesizers. In addition, when used as part of a direct digital frequency synthesizer (DDS), the noise power produced by second-order dithering to 5 bits is greater than the quantization noise introduced by 8 bit digital-to-analog conversion, regardless of sampling rate. As a result, second-order dithering down to 8 bits might be used in a high-performance DDS system, resulting in phase-spur performance that vastly exceeds our example requirement, yet giving a table size half that of the first-order dithered system that just meets the requirement.

Because of the significant noise penalty, second-order dithering down to the minimum number of bits required is considered useful mainly for high-rate applications in which coarse-resolution digital-to-analog converters, e.g., less than 5 bits, are used, or to reduce complexity in those applications where noise power spectral density requirements on the NCO are 100SC.

In applications where one can tolerate the level of noise introduced by using second-order phase dithering, it makes sense to amplitude dither the resulting sinusoid values to reduce the word length to be equal to the reduced phase word length. The amplitude dithering technique is the straightforward addition of a uniformly distributed $[0, A)$ variate, as described in [4]. This will add a white noise component of variance $\frac{A^2}{6}$, which reduces the signal-to-noise ratio by only 0.07 dB. For the -90 dBc example given, a reduction from 16 amplitude bits to 5 amplitude bits would save 59% of the complexity of data path and input/output hardware; if the sinusoid is multiplied digitally, savings in multiplier complexity would be much greater.

5 Practical Dither Generation

In practical systems, a pseudorandom number sequence is substituted for the i.i.d. dither sequence used in the analysis. One class of sequence generators of particular interest are the LFSR pseudo-noise (PN) sequence generators. These generators are easily implemented, requiring only a shift register and a few XOR gates, and yet they generate bit streams with quite good randomness relative to Golumb's randomness postulates [6]. However, the LFSR PN generators produce periodic dither sequences, which makes them not strictly i.i.d. For proper randomness, the generators of each uniform deviate added into the higher-order phase dither variate need to be initialized to different states.

As one might expect, the use of a periodic dither sequence, random within a period, produces a periodic term in the quantization error. This, in turn, will introduce spurious content to the digital sinusoid generated. Such a random, periodic dither sequence with period L , has a power spectrum consisting of a series of L impulses, evenly spaced in frequency, each having equal power. Using the total noise power derived earlier, the resulting spur powers are each less than or equal to $\frac{(M+1)\pi^2\Delta^2}{3L}$, where M is the order of the dither. Thus, for a specified spurious level and dither order, the minimum required LFSR length can be determined. Since the period of the sequence, L , is

exponentially related to the LFSR length, l , $l = 2^l - 1$, designs with reasonable length LFSR PN generators, ($l = 20-30$), will suffice even for very high spectral purity applications.

6 Design Example

The block diagram of a digital sinusoid generator incorporating second-order dithering is shown in Figure 1. The dramatic reduction in look-up table size and arithmetic complexity enabled by the use of second-order dithering makes this -90 dBc spur-free NCO easily implementable in VLSI. The second-order phase dither produces a digital sinusoid contaminated with periodic noise, white to a level of -90 dBc, at a signal-to-noise ratio of 20 dB. Such a digital sinusoid generator is useful as a local oscillator in high-rate (> 100 MHz) applications or in noisy signal environments.

The system parameters are as follows:

Phase bits are in unsigned fractional cycle representation with:

phase accumulator word-length determined by frequency resolution, and

≥ 16 bits prior to addition of the sum of 2 uniform phase dither variates, with ≥ 5 bits after dither addition and truncation;

Amplitude look-up-table with:

≥ 8 entries (using quadrant symmetries) of ≥ 16 bits each normalized so that the sinusoid amplitude equals 2048 16-bit quantization steps less than the full-scale value;

Linear feedback shift register PN generator with ≥ 18 lags producing one 11-bit amplitude dither variate, and

Two LFSR PN generators each with ≥ 24 lags for generation of two 12-bit phase dither variates, initialized to different states.

A 12-bit full adder to sum the two phase-dither variates into a 13-bit result.

7 Simulation Results

Figures 2 through 4 illustrate the spurious improvement of higher-order phase dithering when the look-up phase word-length is constrained to be 5 bits. The frequency is the same for all three figures. The dither signals were generated by a 24-tap LFSR PN generator, forming multi-bit words from the instantaneous contents of the shift register. As such, the dither sequences are not quite white, with a slight lowpass coloration with less than 5 dB drop across the entire band. These simulation results were generated with 512k point, unwindowed FFTs. For display purposes, the positive frequency half of the resulting spectra were compressed by taking the maximum of each block of 64 adjacent power values.

Figure 2 illustrates the worst-case phase spur performance, produced as in [1], at a level of -26 dBc for five-bit phase lookup. Note that in this case, all of

the phase quantization error power is in discrete frequency components, causing the spur and some slight distortion of the fundamental.

Figure 3 illustrates the improvement gained with first-order dithering. The measured average noise power spectral density, before compression by maximum, is -77 dBc per FFT bin, giving a total noise power of -20 dBc. The worst-case phase spur is noticeable at a level of -52.3 dBc, in agreement with theory.

Figure 4 shows the same frequency with second-order dithering. The average noise power is -74 dBc per FFT bin. As can be seen, the spur is substantially reduced to below the level of the noise. In fact, for this frequency, the worst-case, for both the unwithered and the first-order dithering, $E\{\epsilon^3[n]\}$ is a constant δ^3 , making any spurs present $O(\Delta^4)$ in amplitude. The dither signals were generated using two 24-tap LFSR PN generators, with the same generator polynomial.

Figure 5 shows a worst-case result for second-order dithering to 4 bits of phase. The dither sequence was generated as in Figure 4, in order to show spurs clearly, a 16 Megapoint FFT was performed in this case. For display purposes, the positive frequency half the resulting spectrum was compressed by taking the maximum of each block of 512 adjacent power values. The results show residual second-order spurs at the predicted levels for 4 bits, approximately -72 dBc.

8 Conclusion

The technique of higher-order phase dithering for reducing the complexity of high-spectral purity digital sinusoid generators has been presented and analyzed. This technique adds M uniform random deviates to reduce the word length used to represent a sinusoid's phase value while introducing periodic components at a level of $-6(M+1)$ dBc per bit, and a white noise floor. While the noise level added in taking full advantage of even second-order dithering ($M=2$) may be too high for low speed, high SNR applications, the whiteness of the added noise makes this technique ideal for use in high speed (> 500 MHz) integrated-circuit sinusoid generators. Linear feedback shift register-based PN generators can be used as a low-complexity dither generator with this technique, providing a straightforward path to VLSI implementation. For further circuitry savings, the resulting sinusoid samples may be amplitude dithered to reduce their word length without significant loss.

Acknowledgements

The research described in this paper was carried out at the Department of Electrical Engineering and at the Jet Propulsion Laboratory of the California Institute of Technology, under a contract with the National Aeronautics and Space Administration. Michael J. Flanagan was supported in part by a National Science Foundation Fellowship.

References

- [1] G.A. Zimmerman and M.J. Flanagan, "Spur-Reduced Numerically-Controlled Oscillator for Digital Receivers," *26th Annual Asilomar Conference on Signals, Systems and Computers*, Pacific

Grove, California, pp. 517-520, IEEE Computer Society Press, October 1992.

- [2] I. Schuchman, "Dither signals and their effects on quantization noise," *IEEE Transactions on Communication Technology*, vol. COM-12, pp. 162-165, December 1964.
- [3] R.M. Gray and T.G. Stockham, Jr., "Dithered Quantizers," *IEEE Transactions on Information Theory*, vol. IT-39, No. 3, pp. 805-812, May 1993.
- [4] M. Flanagan, G. Zimmerman, "Spur-Reduced Digital Sinusoid Synthesis," to appear *IEEE Transactions on Communications*.
- [5] G.A. Zimmerman, M.F. Garyantes, and M.J. Grimm, "A 640-MHz 32-Megachannel Real-Time Polyphase-FFT Spectrum Analyzer," *25th Annual Asilomar Conference on Signals, Systems, and Computers*, Pacific Grove, CA, pp. 106-110, IEEE Computer Society Press, November 1991.
- [6] Tilborg Henk C.A. van, *An Introduction to Cryptography*, Kluwer Academic Publishers, Boston, pp. 20-28, 1988.

Fig. 1: Block diagram of Design Example

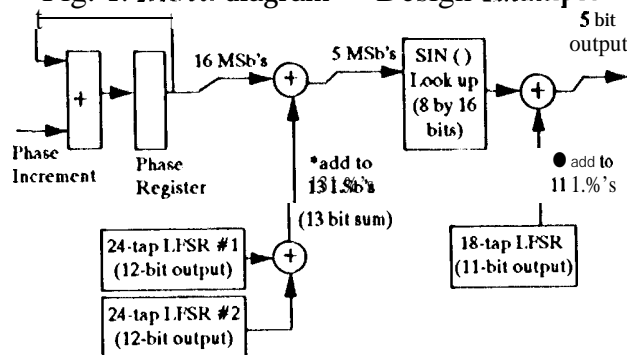


Fig. 2: No Dither, 5 Phase Bits




Fig. 3: 1st Order Dither, 5 Phase Bits




Fig. 4: 2nd Order Dither, 5 Phase Bits




Fig. 5: 2nd Order Dither, 4 Phase Bits

